

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</small> <b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b>					
<b>1. REPORT DATE (DD-MM-YYYY)</b> 05-08-2004		<b>2. REPORT TYPE</b> Final Report		<b>3. DATES COVERED (From – To)</b> 10 February 2003 - 08-Feb-05	
<b>4. TITLE AND SUBTITLE</b>  Evolving Visually Guided Behaviour In 'Hovering' Virtual Agents				<b>5a. CONTRACT NUMBER</b> FA8655-03-1-3033	
				<b>5b. GRANT NUMBER</b>	
				<b>5c. PROGRAM ELEMENT NUMBER</b>	
				<b>5d. PROJECT NUMBER</b>	
<b>6. AUTHOR(S)</b>  Dr. Ray B Lotto				<b>5d. TASK NUMBER</b>	
				<b>5e. WORK UNIT NUMBER</b>	
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Institute of Ophthalmology 11-43 Bath Street London EC1V 9EL United Kingdom				<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>  N/A	
<b>9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b>  EOARD PSC 802 BOX 14 FPO 09499-0014				<b>10. SPONSOR/MONITOR'S ACRONYM(S)</b>	
				<b>11. SPONSOR/MONITOR'S REPORT NUMBER(S)</b> SPC 03-3033	
<b>12. DISTRIBUTION/AVAILABILITY STATEMENT</b>  Approved for public release; distribution is unlimited.					
<b>13. SUPPLEMENTARY NOTES</b>  Final report consists of two documents.					
<b>14. ABSTRACT</b>  This report results from a contract tasking Institute of Ophthalmology as follows: The contractor will develop a series of models that will resolve visual ambiguity in progressively more complex situations. These will begin with the most simple reflectance of objects and, at the end of the grant progress to correlated landscapes with competitive and co-evolving environmental interactions. This will be done through the analysis of neural networks using Bayesian Decision Theory.					
<b>15. SUBJECT TERMS</b> EOARD, Vision, Mathematical Modeling, Neural Networks					
<b>16. SECURITY CLASSIFICATION OF:</b>			<b>17. LIMITATION OF ABSTRACT</b> UL	<b>18, NUMBER OF PAGES</b>  14	<b>19a. NAME OF RESPONSIBLE PERSON</b> JAMES GREENBERG
<b>a. REPORT</b> UNCLAS	<b>b. ABSTRACT</b> UNCLAS	<b>c. THIS PAGE</b> UNCLAS			<b>19b. TELEPHONE NUMBER</b> <small>(Include area code)</small> +44 (0)20 7514 4922

## **EVOLVING VISUALLY GUIDED BEHAVIOUR IN 'HOVERING' VIRTUAL AGENTS**

PI: Dr. R. Beau Lotto

### **BACKGROUND:**

The primary task of vision is to usefully guide behaviour towards objects in the natural world. The challenge in achieving this aim, however, is that the light stimuli that fall on an eye are infinitely ambiguous. The need to overcoming stimulus ambiguity is essential for evolving useful behaviour in natural and artificial agents. And recent research in humans suggests that natural visual systems overcome this challenge by shaping their functional structure wholly according to the statistics of experience. If correct, then the genesis of robust autonomous agents will require understanding the computational principles that correlate evolved structure with the ecological experience that shaped it. Since the necessary description of natural systems is not possible to achieve this aim, we proposed to:

*Stage I:* ...develop a virtual world in which neural net agents 'hover' above a 2D landscape. Since (i) the behaviour and architecture of the networks, (ii) the processes of adaptation, (iii) the structure of the environment, and (iv) probabilistic relationship between stimuli and behavioural significance will be known and manipulated in an environment that incorporates essential aspects of natural ecologies (*behavioral interaction, co-evolution, and stimulus ambiguity*), this evolutionary model of vision provides detail necessary to elucidate the probabilistic basis of visual behaviour.

*Stage II:* ...evolve virtual agents within Mosaic World, and measure their functional and behavioural responses.

*Stage III:* ...create a general statistical (Bayesian) model that correlates neural network functional structure and behaviour with the statistics of the agent's ecological history.

### **ACHIEVED TO DATE:**

Stage I has been completed, resulting the virtual space we call 'Mosaic World'. Rather than give a detailed description of its many different parameters, four time-lapse videos have been created to complement a more superficial description of this virtual space.

### **BRIEF DESCRIPTION OF COMPLETED STAGE I:**

#### *The Space:*

Mosaic World is composed of an  $n \times n$  array of surfaces (the actual size of the matrix is unlimited). As in nature, each surface in this space is defined mathematically by its 'reflectance efficiency function', which describes the nature of its reflectance between 400 and 700 nm. Though randomly generated, reflectances are constrained by the known

statistics of natural images. Using a novel growth algorithm, an infinite number of highly controllable, though indeterminate worlds can be generated.

#### *Illumination:*

As in nature, surfaces in Mosaic World are lit by extended light sources, whose spectral characteristics change continuously in space and time. The size of each illuminant can range from a single surface to the entire matrix, and can exist for one to an infinite number of time-steps. The overall illumination intensity also varies, thus mimicking natural 'day/night' cycles. As with the surface reflectance matrix, the illumination matrix is created with the same unique growth algorithm.

#### *Agents:*

Each of an agent's behaviours (e.g., movement, eating, and mating) is controlled by an evolvable artificial neural network, the attributes of which (e.g., number and location of input nodes, number and location of the hidden nodes, and size and efficiency of the connection matrix) can be independently evolved. To achieve this, we have designed a unique method of genetic programming that encodes each agent element, not as genes, but as objects. The resulting functional structure of each network is wholly indeterminate.

#### *Task:*

There is only a single selection criterion in Mosaic World: Survival. To survive, agents must obtain resources from the environment, while minimising costs. Resources are obtained either directly from the surfaces (herbivores), or from other agents (carnivores). The value of each is determined directly by the colour of the surface (in the case of herbivores), or by the energy store of the consumed agent (in the case of carnivores). As such, agents in Mosaic World must overcome the inherent ambiguity of the spectral information they receive from their environment to survive.

### BRIEF DESCRIPTION OF COMPLETED STAGE II:

After completing the visual world (Stage I), the aim of Stage II was to determine/demonstrate that it is possible to evolve useful behaviour within this ecologically relevant virtual space. This we have also achieved. Indeed, our initial results are currently under review at GECCO. For sake of clarity, we've also attached that manuscript, and so will not expand on those results here. Nonetheless, in completing this stage, we have completed exactly what we aimed to achieve in the first two years that funding was provided. As the project is being terminated, we will not be able to pursue the final (Stage III), where the aim was to use the model to determine/generate robust statistical algorithms that could be generalized to autonomous visually-guided robotic systems in natural environments.

# Evolving Visually Guided Agents in an Ambiguous Virtual World

Ehud Schlessinger<sup>1</sup>, Peter J. Bentley<sup>2</sup>, and R. Beau Lotto<sup>1</sup>

<sup>1</sup> Institute of Ophthalmology, University College London, 11-43 Bath Street,  
London EC1V 9EL  
{e.schlessinger, lotto}@ucl.ac.uk

<sup>2</sup> Department of Computer Science, University College London, Gower Street,  
London WC1E 6BT  
P.Bentley@cs.ucl.ac.uk

**Abstract.** The fundamental challenge faced by any visual system within natural environments is the ambiguity caused by the fact that light that falls on the system’s sensors conflates multiple attributes of the physical world. Understanding the computational principles by which natural systems overcome this challenge and generate useful behaviour remains the key objective in neuroscience and machine vision research. In this paper we introduce Mosaic World, an artificial life model that maintains the essential characteristics of natural visual ecologies, and which is populated by virtual agents that – through ‘natural’ selection – come to resolve stimulus ambiguity by adapting the functional structure of their visual networks according to the statistical structure of their ecological experience. Mosaic World therefore presents us with an important tool for exploring the computational principles by which vision can overcome stimulus ambiguity and usefully guide behaviour.

## 1 Introduction

While computers can outplay the top chess Grandmasters, they cannot recognise a flower in a meadow with the same proficiency of a bumblebee. The reason is that, unlike computers, the bee’s miniature brain, which is made up of 240 times fewer neurons than humans have retinal receptors, has – like all natural visual systems – adapted to resolve the most fundamental challenge of natural visual ecology, which is that light stimuli are fundamentally ambiguous with respect to their natural provenance.

Though seemingly counterintuitive, since our internal representation of the world seems so **un**ambiguous, consider (i) that the world is composed of many kinds of objects under many different kinds of lights, and (ii) that vision does not have direct access to that world. Rather, the only information that any visual system has access to is the light that falls on its eye/sensors (i.e., stimuli). Now, assuming that useful vision requires generating behaviour according to the sources of stimuli, then we have

what seems an impossible problem. The problem is this. The quality and intensity of any stimulus is determined from both an object's reflectance and illumination *simultaneously*. Since changing either parameter changes the stimulus accordingly, there is no way to deduce from the stimulus itself its underlying source, which could be any number of different reflectance/illumination combinations. Thus, the task of recognising an object under different illuminants is like solving the equation  $x \cdot y = z$  for  $x$  without ever knowing  $y$ . And yet, natural systems – from bees to humans – solve this problem continuously when recognising objects in natural environments. Understanding how this is achieved is essential to understanding natural vision, and for creating successful artificial visual systems.

While the physiological mechanisms are currently unknown, recent neuroscience research suggests natural systems overcome the problem of stimulus ambiguity by generating visual behaviour probabilistically, by shaping the functional structure of visual processing empirically in ontogeny and phylogeny (Lotto and Purves, 1999, 2000, 2001; Purves and Lotto, 2003). The premise of this hypothesis is that to be useful, vision need not accurately represent the world, but generate statistical constructs based on what proved useful in past experience. If correct, then to understand vision will require understanding in quantitative terms of the statistical relationship between images, their *behavioural* consequences in past experience, and the functional structure of adaptive network systems, information which is unknown, if not unknowable in natural systems. It is within this context that Artificial Life presents us with a necessary tool for prosecuting the empirical basis of vision.

Indeed, researchers in many fields have used A-Life models to explain biological phenomena, such as – but not limited to – ecological dynamics of growing plants (Dyer et al, 2001; Dyer and Bentley, 2002), sexual selection (Collins and Jefferson, 1992), foraging behaviour of Anolis lizards (Koza et al, 1992), and even war and peace between virtual nations (Unemi et al, 2003). Within the context of vision research, Liese et al (2001) created an A-Life simulation in which virtual light sensors evolved according to whether or not they obtained sufficient light from their environment. In another model, Kortmann et al (2001) evolved a population of visuo-motor systems to investigate the trade-off between spatial and temporal resolution that occur in biological systems. Finally, Terzopoulos and Rabie (1995) created a realistic simulation of a population of artificial fish to facilitate "...active vision research".

While the above visual A-Life models were useful in their own right, they did not, nor could they help us to address questions about how natural vision systems deal with stimulus ambiguity. Mosaic World was therefore created to overcome this obstacle by preserving these essential characteristics of natural visual ecology.

## 2 System

Mosaic World is an A-Life, virtual environment made up of 2D 'coloured' surfaces under simple-to-complex patterns of 'coloured' illumination. The virtual agents, the 'critters', that inhabit Mosaic World have one task: to survive, which they do by eating good resources and avoiding bad ones. Furthermore, as a stable population re-

quires new members, the agents multiply through sexual and asexual reproduction. The behaviour of each agent is controlled by an artificial 3D neural network. All attributes of the networks – from receptor number and spectral sensitivity, to hidden node number, to the connections between the nodes – are evolvable.

## 2.1 Environment

The ground matrix of Mosaic World is made up of 100x100 ‘surfaces’. Beyond the edges of the world, there is nothing: A critter attempting to travel more than one surface beyond the edge dies. Within the surface matrix, there are holes. Entering a hole has the same effect as stepping over the world’s edges.

Every surface is defined in terms of its location and reflectance function, which describes its proficiency at reflecting wavelengths within the range of visible light (400 to 700 nm). Because Critters ‘eat’ surfaces, or more specifically their constituent wavelengths, a surface’s reflectance function determines its resource value, and hence its behavioural relevance: A surface that reflects wavelengths that the experimenter decides are ‘good’ will add positively to the agent’s resource intake, whereas eating from surfaces that reflect only ‘bad’ wavelengths will add negatively to the agent’s resource intake. As a surface is ‘consumed’, it becomes increasingly less reflective, and hence more transparent, eventually becoming invisible. Through time it will regenerate its reflectance at a rate pre-determined by the experimenter.

In nature, because objects and illuminants are extended in space, two neighbouring points in any image are more likely to have the same physical characteristics than are two points further away from each other. To model this statistical relationship in Mosaic World, a CA (cellular automata) algorithm was created that enables complete control over such clustering across the ground matrix (see Figure 1 for algorithm description), as well as the relative proportion of negative-to-positive wavelength resources in the world. Thus, the surface matrix can be customised in terms of reflectance/illumination size, which can vary from one pixel to the entire world, as well as each clusters’ shading, which can be gradual changes to random transitions.

Illumination of Mosaic World is controlled by simulated external light sources, each modelled as a discrete wave function in the range between 400nm and 700nm. The spatial distribution of illumination across the surface matrix is controlled using the same CA algorithm used to create the surface matrix. Illumination can also change in time, both in terms of quality (i.e., its spectral distribution), size and movement across the space, all at rates pre-defined by the researcher. See Figure 2 for a few examples.

It’s important to stress that, as in natural environments, each stimulus arising from each surface in Mosaic World is determined by the relative contribution of its reflectance and its illumination: *Stimulus (S) = Reflectance (R) • Illumination (I)*. **It is this ambiguous stimulus that is presented to the virtual agents**, with the consequence that there is no direct way for the agent’s sensors to estimate a given surface’s type from the stimulus alone.


Constants:

- CLUST\_PARAM = [0,1], the probability of new surfaces being identical to neighbours
- SHADE\_PARAM = [0,1], similarity factor of non-identical surfaces to neighbours.
- POS\_RES\_PARAM = [0,1], the probability of generating a positive surface/resource
- NEG\_RES\_PARAM = [0,1], the probability of generating a negative surface/resource

Method of operation:

- (a) An empty world without any surfaces is created.
- (b) A predefined number of surface seeds are generated. Every seed is placed in a random location in the world and is initialised with a reflectance using POS\_RES\_PARAM and NEG\_RES\_PARAM (or alternatively, a random reflectance is initialised).
- (c) New surfaces are being 'thrown' into the world at random locations. If there are no adjacent surfaces near a new surface, it is removed from the world.
- (d) If there are one or more adjacent surfaces, the new surface gets a reflectance influenced by one of its neighbours, randomly chosen.
  - (i) Using CLUST\_PARAM, the new surface is deemed to be identical or not to its neighbours. E.g. A value of 0 determines that there's 0% it'll be identical to one of its neighbours, and 100% it'll be non-identical.
  - (ii) If the new surface is deemed to be non-identical to its neighbour, using SHADE\_PARAM it is determined how similar it would be: A value of 0 determines that every wavelength in the reflectance function will be randomly changed by 0 (so it'll be identical to its neighbouring surface), a value of 1 determines that every wavelength in the reflectance function will be randomly changed by -1 to +1 (so it'll probably be very different from its neighbour).
- (e) This process is iterated until all the surface positions have been created.
- (f) Measure world statistics: unless within defined ranges, go back to step (a).

Figure 1: The Shading/Clustering algorithm



**Figure 2: A few sample worlds. [a] Large clusters, uniform ‘white’ illumination. [b] Same world as [a] with multiple illuminants. [c] Many small clusters, uniform ‘white’ illumination, 3 holes. [d] Same world as [c] with multiple illuminants.**

## 2.2 Critters

‘Critters’ are the inhabitants of Mosaic World. Every critter has *field of view* (the area it receives stimuli from the environment), and an *orientation* (the direction it is facing). Upon instantiation, all critters are given a certain amount of energy, which decreases in time. The amount of energy lost in each time-step is determined by the following: rate of motion (faster speeds cost more), turning (in increments of 90 degrees), eating (big ‘bites’ cost more: bite size is ‘chosen’ by the critter itself), and reproduction. Also, the larger a critter’s neural network structure, the more costly is each time-step. If a critter’s energy level drops to 0, it dies, giving it a strong incentive to gain energy by consuming surface reflectances (as described above), or by preying on other critters (even though all critters are instantiated as omnivores, they can become exclusively herbivores or carnivores in time).



Critter population is maintained in Mosaic World through sexual and asexual reproduction. To facilitate sexual reproduction, critters are also instantiated with a transmittant colour, which can evolve, and can thus be used as a cue for recognising conspecifics (as well as predators and prey). Depending on the mode of reproduction, an offspring's genome will either be a mutated version of its parent (asexual reproduction), or will be a mutated and recombined version of its parents' genomes (sexual reproduction). The offspring's starting energy is transferred from its parent/s (thus, reproduction costs energy).

Each critter's phenotype is determined by its genome, which is a conceptual entity, that defines all the traits of a critter: Visual system (all details regarding the receptors: position, tuning, peak, activity), 3D Neural network (number and position of hidden units, definition of connection weights: value, starting unit, ending unit or coordinate), transmittance (the critter's colour). There is a one to one relationship between the genes in the genome and a critter's phenotype.

Each physical attribute of the critter is treated as an object, and each object can either mutate, or parts of it can be recombined with comparable objects from other critters during sexual reproduction. Thus, the only time the genome is explicitly displayed is when a critter is saved for analysis and is stored in a text file (see summarised genome in Figure 3). Of course the most important aspect of each critter's phenotype is its 'visual brain', which translates the ambiguous stimulus information into useful behaviour.

```
Size: 1x1
Transmittance wave function: 0.33, 0.35, 0.37, 0.40, 0.42, 0.44, 0.45, 0.46, 0.46,
0.47, 0.48, 0.50, 0.52, 0.54, 0.56, 0.58, 0.57, 0.5, 0.53, 0.51, 0.49, 0.51, 0.52, 0.54,
0.55, 0.57, 0.58, 0.58, 0.59, 0.60, 0.61
3D Neural network (partially connected):
Visual layer: 3 units:
- Health unit
- Receptor 1: coordinate: [0,-1], peak: 680nm, tuning: 0.01226, active.
- Receptor 2: coordinate: [0,0], peak: 400nm, tuning: 0.02868, active.
Hidden layer: 4 units:
- Hidden unit 1: coordinate [-1,-1]
- Hidden unit 2: coordinate [0,0]
- Hidden unit 3: coordinate [2,0]
- Hidden unit 4: coordinate [-1,1]
Output layer: 7 units
Active Connections: 33
Unconnected connections: 1
```

**Figure 3: Sample summarised critter genome**

### 2.2.1 Visual Brain


An agent's visual brain comprises a modified 3D feedforward neural network. The 3D network is composed of multiple 2D layers (see Figure 4 for illustration). The

input layer is the *visual layer*, which contains receptors (these are essentially modified input units) and a *health monitor unit*, which receives the percentage of the critter's remaining health. The hidden layers contain standard hidden units. The output layer contains output units, which determine the critter's behaviour: turn left or right, move forward or stay in the same position, sexually reproduce with a critter standing where it is, asexually reproduce, bite surface and bite a critter standing where it is. Every unit in the network has an [x,y] coordinate relative to the critter's centre, which defines its location in the layer it is placed in. Thus, an important and novel attribute of our design is that networks of vastly different architectures can be crossed over during sexual reproduction, as each network possesses the same virtual body/coordinate reference frame.

The receptor's position in the input layer also determines the location – relative to the critter's centre – where it detects light from Mosaic World (i.e., its visual 'receptive field'). How it responds to that light is determined by its spectral sensitivity function, as in natural systems, between 400nm-700nm. The receptor's *peak* sensitivity can be for any wavelength within this range (at increments of 10nm). The receptor's *tuning* defines its half bandwidth (i.e., the number of wavelengths to which it will respond). A receptor can be either active or inactive. Inactive receptors do not participate in subsequent processing, but are nonetheless inherited by offspring; these are discarded if inactive for a long period of evolutionary time.

Every receptor has an adaptation factor, which is determined from one of the network's output neurons, thus enabling the critter to adapt to different ranges of stimuli. A similar mechanism is employed in nature and enables humans to see both at night and during the day.

The units of the network communicate through connection weights that can extend between units from higher layers to lower layers, and can also connect units to empty coordinates in the network ('partial' connections). Connections can be active, inactive, or completely nonexistent. Inactive, 'partial' or 'missing' connections do not participate in the feed-forward process. Connections are discarded if inactive for long periods of evolutionary time.



**Figure 4: ‘Visual Brain’ - 3D neural network. This brain has three layers (one hidden layer). The visual layer contains three receptors (one highly tuned, the other two possess lower tuning values). The visual layer is connected to the hidden layer, specifically to five hidden units and two empty coordinates (partial connections). The hidden layer is connected to the output layer.**

### 2.3 Evolution

To maintain an open-ended system, Mosaic World’s evolution utilises a genetic algorithm with no fixed population size and no explicit fitness function. The critters themselves decide when to reproduce (sexually or asexually) by activating the appropriate output neurons. Critters survive if they can eat good resources, avoid predators and stay on the world. Thus, there’s an implicit evolutionary selection pressure to improve all traits that increase such skills. For the population to survive, the virtual agents must balance reproduction with resource consumption. Otherwise, there may be too many critters for the world to sustain, or too few critters to maintain the population.

Because at time-step 0 all critters are randomly instantiated, a statistical consequence of this is that the initial population sometimes dies. When this happens, a new population of random critters is instantiated, with the caveat that 20% are mutated clones of critters that showed general promising surviving skills (a combination of survival age and mating amount).

Crossover takes place during sexual reproduction. During crossover, a random point is selected on each network layer of both mating critters. All 3D layers of each critter brain are ‘sliced’ at this point, and recombined with the half from the other

critter in creating offspring. Not only does this process create novel patterns of node distributions, but it also creates connections between nodes that existed in neither parent. For instance, if a ‘partial’ connection is obtained from one parent, and a hidden node at the corresponding coordinate that previously lacked a connection, from the other parent, then this becomes a novel connection that didn’t exist in either parent previously. Crossover also combines the transmittance functions of both mating critters.

Mutation takes place during sexual and asexual reproduction. Mutation in a genome occurs through several different ways: 1. *Drift mutation* causes a receptor to randomly switch positions in the visual layer (0.3% per receptor). 2. *Peak mutation* changes the receptor’s peak (2% per receptor). 3. *Tuning mutation* changes the value of the critter tuning (2% per receptor, changes by up to 0.2). 4. *Activation state mutation* changes the state (active or inactive) of receptors and connections affected (0.3% per receptor and connection). 5. *Structural Mutations* that change the brain structure, such as 5a. *Delete Receptor* (0.3% per receptor), 5b. *Delete Hidden Unit* (deleting a hidden unit leaves unconnected connections; 0.3% per unit), 5c. *Delete Connection* (0.1% per connection), 5d. *Add Connection* (1%), 5e. *Add Receptor* (1%) and *Add Hidden* (1%). 6. *Value mutations* affect the value of the connection weights (changes by up to 1). And 7. *Transmittance mutations* affect the transmittance defining the colour of a critter (per wavelength, by up to 0.05).

### 3 Experiment

A principle hypothesis of Mosaic World is that critters that evolve within relatively ‘simple’ worlds composed of one uniform illuminant, where there is a simple one-to-one relationship between stimuli and their source, will survive less well in **ambiguous** environments, under multiple illuminants, where there is a many-to-one relationship between a stimulus and its source. In contrast critters that evolved within such complex worlds will survive well in these worlds. To test this, agents were presented with a surface matrix under one illuminant, or under multiple illuminants. To survive critters would have to learn how and when to move, when to mate, how to avoid the world’s edge and its holes, as well as recognise positive resources and avoid bad ones. In this experiment, short wavelengths were ‘good’ (increased the critter’s health), and long wavelengths were bad (decreased the critter’s health). Thus, agents needed to recognise surfaces that reflected predominantly short wavelengths (which would appear ‘bluish’ to the human visual system) under multiple lights, and avoid those that would appear ‘reddish’ to human observers. Green, grey and purple surfaces, and every combination, as they reflect both short and long wavelengths to the same extent, would – on balance – offer no reward (as they would add as much to the health of the agent as they take away). Of course determining the identity of good and bad surfaces becomes increasingly difficult as the number of illuminants in the world increases, which increases the number of stimuli arising from each surface accordingly.

Twelve randomly generated worlds were created using the identical methods and surface/illuminant statistics. Critters were evolved on all 12 worlds under the two

different conditions: (1) constant illumination (a uniform illuminant with an intensity of 0.6 at each wavelength), and (2) multiple illuminants that change in location and quality every 50 time steps. After 550,000 time steps, all 24 runs were stopped, and the 5 oldest critters were taken from each run. This criterion for selection is somewhat arbitrary, but there is no easy way to measure the fitness of evolved critters, and survival age seems to be positively correlated with fitness.

A test world was created, and 3 copies of each of the 5 chosen critters were placed in it (for a total of 15 critters per test). The test world was run under uniform illumination, and under multiple illuminants. Every such run was performed 5 times, for a total of 120 runs. The comparison of the critters was a simple survival test: the average survival time of the critters in the test world was calculated and compared. Critters that survived till the end of the test run (10,000 time steps) were assumed to have died then.

## 4 Results

Condition of Evolution	Average survival under a uniform illuminant	Average survival under multiple illuminants
Random critters	86.66	88.34
Evolution under a uniform illuminant	3067.72	1252.8
Evolution under multiple illuminants	3520.43	1573.46

**Table 1: Average survival times for critters in test worlds; broken down according to evolution with uniform or multiple illuminants.**

Number of Visual Receptors	Survival with uniform illuminant	Survival with multiple illuminants
1	3185.55	1302.48
2	3989.25	1952.83

**Table 2: Average survival time for critters that evolved under multiple illuminants; broken down according to critters that evolved one and two receptors.**

Table 1 shows the resulting data of the average survival time of critters that evolved in multiple or uniform illuminant worlds tested under multiple and uniform illuminant test worlds. As can be seen in the first row, randomly instantiated critters (the control group) did not survive for long in either world. When comparing the second and third rows, it is also clear that critters that evolved under multiple illuminants survived longer than those that evolved under a uniform illuminant when tested in either test world. Thus, the critters with the best performance in all test worlds were critters that evolved under multiple illuminants.

It is interesting to note that when evolved under uniform illumination, only 1 out of 12 test critters had more than one visual receptor. This was not the case, however, under multiple illuminant runs, where 5 out of the 12 test critters evolved multiple receptors. We therefore tested whether there is a behaviour value in adapting multiple receptor types by comparing the survival rates for critters with one or two receptors, which is shown in Table 2. As can be seen in column 1 of Table 2, the survival rate for single receptor critters that evolved in multiple illuminants worlds was very similar to the survival rate of critters that evolved under a single illuminant (and that also evolved a single receptor type). In contrast, critters that evolved two receptor types survived much longer on average - under both conditions, as is shown in column 2. Thus, adapting multiple receptor types, which is fundamental for colour vision, is important in more complex conditions of illumination, presumably because it enables critters to better recognise surfaces under more varied conditions.

## 5 Conclusions

The above experiments confirm the hypothesis that ambiguous environments provide a unique challenge that requires specific adaptations in order to survive and thrive. Once obtained, individuals possessing it in general are more adapted for ambiguous and non-ambiguous environments alike.

Since this challenge is universal in natural environments, it is expected that using the Mosaic World framework experiments examining multiple aspects of natural vision can be run. By creating conditions inspired by their natural counterpart and seeing what strategies the evolved critters use to survive, it is anticipated that exciting and important insights into biological vision can be gained. Furthermore, using reverse engineering of the evolved critters, simple and efficient algorithms that can benefit all vision research may be derived.

## References

1. Collins R. J. and Jefferson D. R. (1992) The evolution of sexual selection and female choice. In F. J. Varela and P. Bourgine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 327--336, Cambridge, MA, 1992.
2. Dyer, J. and Bentley, P. J. (2002) PLANTWORLD: Population Dynamics in Contrasting Environments. A late-breaking paper in *Proc. of the Genetic and Evolutionary Computation Conference (GECCO) 2002*.
3. Dyer, J. R., Bentley, P. J. and Shah, P. (2001) PLANTWORLD: The Evolution of Plant Dormancy in Contrasting Environments. A late-breaking paper In the *Genetic and Evolutionary Computation Conference (GECCO) 2001*. pp. 67-74.
4. Kortmann, R., Postma, E. and van den Herik, J. (2001) Evolution of visual resolution constrained by a trade-off. *Artif Life*. 2001 Spring; 7(2):125-45.
5. Koza, J. R., Rice, J. P., and Roughgarden, J. (1992) Evolution of food foraging strategies for the Caribbean Anolis lizard using genetic programming. *Adaptive Behavior*. Volume 1, number 2, pages 47-74.

6. Liese, A., Polani, D., and Uthmann, T. A (2001) study of the simulated evolution of the spectral sensitivity of visual agent receptors. *Artif Life*. 2001 Spring; 7(2):99-124.
7. Lotto, R.B. and Purves, D. (1999) From the Cover: The effects of color on brightness. *Nature Neuroscience* 2:1010-1014.
8. Lotto, R.B. and Purves, D. (2000) From the cover: An empirical explanation of color contrast. *Proceedings of the National Academy of Science USA* 97:12834-12839.
9. Lotto, R.B. and Purves, D. (2001) From the cover: An empirical explanation of the Chubb illusion. *Journal of Cognitive Neuroscience* 13:547-555.
10. Purves, D.P. and Lotto, R.B. (2003) *Why we see what we do: A wholly probabilistic strategy of vision*. Sinaur Associates INC. (Sunderland Massachusetts) and Macmillan Press (London, UK).
11. Terzopoulos D. and Rabie T. (1995), "Animat vision: Active vision in artificial animals," in *Int. Conf. on Computer Vision*, 1995, pp. 801--808.
12. Unemi T., Kaneko Y. and Takahashi I. (2003), War and Peace among Artificial Nations - a model and simulation based on a two-layered multi-agent system, *ECAL 2003*, Springer-Verlag LNCS/LNAI series, LNAI 2801, pp. 146-153, 2003.